
SHUANGJIAN ZHANG, CNRS, ENS Paris

Wasserstein Control of Mirror Langevin Monte Carlo

Discretized Langevin diffusions are efficient Monte Carlo methods for sampling from high dimensional target densities that are log-Lipschitz-smooth and (strongly) log-concave. In particular, the Euclidean Langevin Monte Carlo sampling algorithm has received much attention lately, leading to a detailed understanding of its non-asymptotic convergence properties and of the role that smoothness and log-concavity play in the convergence rate. Distributions that do not possess these regularity properties can be addressed by considering a Riemannian Langevin diffusion with a metric capturing the local geometry of the log-density. However, the Monte Carlo algorithms derived from discretizations of such Riemannian Langevin diffusions are notoriously difficult to analyze.

In this talk, we consider Langevin diffusions on a Hessian-type manifold and study a discretization that is closely related to the mirror-descent scheme. We establish for the first time a non-asymptotic upper-bound on the sampling error of the resulting Hessian Riemannian Langevin Monte Carlo algorithm. This bound is measured according to a Wasserstein distance induced by a Riemannian metric ground cost capturing the squared Hessian structure and closely related to a self-concordance-like condition. The upper-bound implies, for instance, that the iterates contract toward a Wasserstein ball around the target density whose radius is made explicit. Our theory recovers existing Euclidean results and can cope with a wide variety of Hessian metrics related to highly non-flat geometries. This talk represents joint work with Gabriel Peyré, Jalal Fadili, and Marcelo Pereyra.